arXiv:1412.5995v3 [q-bio.QM] 30 Jun 2015

Fast and accurate approximate inference of transcript expression from RNA-seq data

James Hensman*‡[1], Panagiotis Papastamoulis†‡[2], Peter Glaus[3], Antti Honkela[4] and Magnus Rattray[2]

[1]Sheffield Institute for Translational Neuroscience (SITraN), UK
[2]Faculty of Life Science, The University of Manchester, UK
[3]School of Computer Science, The University of Manchester, UK
[4]Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Finland

July 1, 2015

Abstract

Motivation: Assigning RNA-seq reads to their transcript of origin is a fundamental task in transcript expression estimation. Where ambiguities in assignments exist due to transcripts sharing sequence, e.g. alternative isoforms or alleles, the problem can be solved through probabilistic inference. Bayesian methods have been shown to provide accurate transcript abundance estimates compared to competing methods. However, exact Bayesian inference is intractable and approximate methods such as Markov chain Monte Carlo (MCMC) and Variational Bayes (VB) are typically used. While providing a high degree of accuracy and modelling flexibility, standard implementations can be prohibitively slow for large datasets and complex transcriptome annotations.

Results: We propose a novel approximate inference scheme based on VB and apply it to an existing model of transcript expression inference from RNA-seq data. Recent advances in VB algorithmics are used to improve the convergence of the algorithm beyond the standard Variational Bayes Expectation Maximisation (VBEM) algorithm. We apply our algorithm to simulated and biological datasets, demonstrating a significant increase in speed with only very small loss in accuracy of expression level estimation. We carry out a comparative study against seven popular alternative methods and demonstrate that our new algorithm provides excellent accuracy and inter-replicate consistency while remaining competitive in computation time.

Availability: The methods were implemented in R and C++, and are available as part of the BitSeq project at `github.com/BitSeq`.

*james.hensman@sheffield.ac.uk
†panagiotis.papastamoulis@manchester.ac.uk
‡These two authors contributed equally.

The method is also available through the BitSeq Bioconductor package. The source code to reproduce all simulation results can be accessed via `github.com/BitSeq/BitSeqVB_benchmarking`
Keywords: RNA-Seq, Transcript expression estimation, Bayesian inference, Variational Bayes.

1 Introduction

RNA-seq is a technology with the potential to identify and quantify all mRNA transcripts in a biological sample (Mortazavi et al., 2008). Some of these transcripts come from different isoforms or alleles of the same genes or from closely related homologous genes, and consequently they may share much of their primary sequence. Currently popular RNA-seq technologies generate short reads that must be aligned to the genome or transcriptome in order to quantify expression levels. In some cases the observed reads could originate from several different transcripts and there may be few reads that are useful to distinguish these transcripts. It is therefore a challenging statistical problem to uncover the expression levels of closely related transcripts. A recent assessment confirms this by showing significant variability between results obtained using different computational pipelines (SEQC/MAQC-III Consortium, 2014).

Probabilistic latent variable models, in particular mixture models (Jiang and Wong, 2009; Li et al., 2010; Katz et al., 2010; Li and Dewey, 2011; Turro et al., 2011; Glaus et al., 2012; Trapnell et al., 2013; Nariai et al., 2013) provide a popular and effective approach for inferring transcript expression levels from RNA-seq data. Such models can be used to deconvolve the signal in the read data, assigning reads to alternative, pre-defined transcripts according to their probability of originating from each. The term mixture model derives from the interpretation of the data as being derived from a mixture of different transcripts, the mixture components, with each read originating from one component. Although reads originate from only one component they may map to multiple related components, resulting in some ambiguity in their assignment. Transcript expression levels are model parameters (mixture component proportions) that have to be inferred from the mapped read data. Due to their probabilistic nature these models can fully account for multiple mapping reads, complex biases in the sequence data, sequencing errors, alignment quality scores and prior information on the insert length in paired-end reads. Mixture models have been successfully applied to infer the proportion of different gene isoforms or allelic variants in a particular sample (Jiang and Wong, 2009; Katz et al., 2010; Turro et al., 2011), for inferring gene and isoform expression levels (Li et al., 2010; Li and Dewey, 2011; Roberts and Pachter, 2013; Trapnell et al., 2013; Mortazavi et al., 2008) and for transcript-level differential expression calling (Glaus et al., 2012; Trapnell et al., 2013).

Inference in latent variable models such as these can be carried out by Maximum Likelihood (ML) or Bayesian parameter estimation. In ML the choice of parameters that maximises the data likelihood is obtained through a numerical optimisation procedure. In the case of mixture models a popular choice of

algorithm is the Expectation Maximisation (EM) algorithm, as first applied to this model and expressed sequence tag data by Xing et al. (2006) and later to RNA-seq data by Li et al. (2010). For Bayesian inference the most popular approach is Markov chain Monte Carlo (MCMC) and for the case of mixture models a Gibbs sampler is most often used (Katz et al., 2010; Li and Dewey, 2011; Glaus et al., 2012). An advantage of Bayesian inference is that one obtains a posterior probability over the model parameters rather than just a point estimate. This provides a level of uncertainty in the inferred transcript expression levels as well as information about the covariation between estimates for closely related transcripts. The uncertainty information can be usefully propagated into downstream analysis of the data, e.g. calling differentially expressed transcripts from replicated experiments (Glaus et al., 2012).

A Bayesian method, BitSeq, was proposed in which inference was carried out using a collapsed Gibbs sampler (Glaus et al., 2012). The method was shown to perform well, especially for the task of inferring the relative expression of different gene isoforms and for ranking transcripts according to their probability of being differentially expressed between conditions. However, for typical modern RNA-seq datasets with hundreds of millions of read-pairs the Gibbs sampler can be inconveniently slow, creating a computational bottleneck in applying a Bayesian approach. As the volume of data continues to grow and gene models are becoming more complex as more alternative transcripts are discovered, more efficient inference algorithms are required so that Bayesian methods can be used to provide practical computational tools.

An alternative approach to Bayesian inference is to use deterministic approximate inference algorithms such as Variational Bayes (VB) (reviewed in Bishop, 2006). While MCMC algorithms are attractive due to their asymptotic approximation guarantees, VB often provides a much faster method to obtain a good approximation to the posterior distribution. For models where Gibbs sampling can be applied there is typically a closely related VB Expectation Maximisation (VBEM) algorithm. In this contribution we show how VB can be used to massively speed up inference in the BitSeq model for transcript expression-level inference. We show that the mean transcript expression level estimates are very close to those obtained with MCMC. We use a recent formulation of VB (Hensman et al., 2012) which is shown to provide a greater speed up when compared to a more standard VBEM algorithm. Our new algorithm is implemented in the most recent version of the BitSeq, allowing the method to be applied to much larger RNA-seq datasets in equal computing time.

An alternative VB method, TIGAR, was recently proposed for the same problem using a standard VBEM algorithm (Nariai et al., 2013). The assumptions made in our approximation are similar to those used in TIGAR, but the empirical comparisons herein show that our proposed method performs better in terms of computation time and required memory, while also providing improved accuracy on real and simulated data. The improvement in terms of reduced computational cost is due to our adoption of a novel VB method. Furthermore we investigate the effects of the variational assumption in this problem, and compare empirically to results using the gold standard, MCMC.

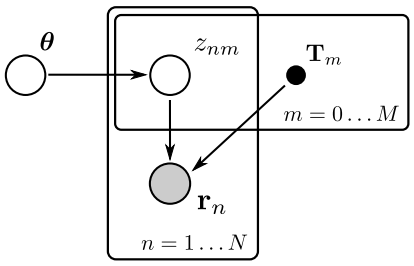

Figure 1: Graphical model of the RNA-seq mixture problem. Given a known Transcriptome $\mathbf{T}$ and some observed reads $\mathbf{R}$, the inference problem is for $\boldsymbol{\theta}$ through the latent variables $\mathbf{Z}$.

The paper is organised as follows. In Section **??** we review the original BitSeq probabilistic model and describe our new inference algorithm, BitSeqVB, explaining the principles underlying our improved optimisation scheme. In Section 3 we benchmark our new method against the original BitSeq algorithm and six popular alternative methods using realistic simulated data and real human RNA-Seq data. We consider accuracy in terms of expression estimation, relative with-gene transcript proportions and between-replicate consistency. We also compare the computation time required for all methods and compare the new VB algorithm to more standard MCMC and VBEM inference algorithms.

2 Methods

Our probabilistic model of RNA-seq follows Stage 1 of Glaus et al. (2012), and is similar to that used by RSEM. We summarise our notation in Table 1. The probabilistic model is shown using standard directed graphical notation in Figure 1. Here we have focussed on the mixture part of the analysis, assuming that the model which associates reads to transcripts (i.e. $p(\mathbf{r}_n \mid \mathbf{T}_m)$) is known. Following BitSeq (Glaus et al., 2012), we compute this part of the model a priori, with parameters estimated from uniquely aligned reads. We consider RNA-seq assays independently, computing an approximate posterior for the transcript proportions $\boldsymbol{\theta}$ in each assay. Subsequent analysis such as differential expression can be done using the estimated distributions of each assay.

2.1 The generative model

Transcript fragment proportions

The generative model for an RNA-seq assay is as follows. We assume that the experiment produces of collection of RNA fragments, where the abundance of fragments derived from transcript T_m in the assay is θ_m. Fragments are then sequenced in these proportions, so that the prior probability of any fragment corresponding to transcript T_m is θ_m. Introducing a convenient allocation vector

Table 1: Summary of notation

N	Number of reads in the dataset.
M	Number of transcripts in the transcriptome.
$\mathbf{r}_n$	The n^{th} read.
$\mathbf{R}$	The collection of reads.
$\mathbf{T}$	The transcriptome.
T_m	The m^{th} transcript.
θ_m	Proportion of transcript T_m in the sample.
z_{nm}	Binary: $z_{nm} = 1$ if read n comes from transcript m.
$\mathbf{z}_n$	Allocation vector of the n^{th} read.
$\mathbf{Z}$	Collection of all allocation vectors.
ϕ_{nm}	Approximate posterior probability of $z_{nm} = 1$.
γ_{nm}	Re-parameterisation of ϕ_{nm}.

$\mathbf{z}_n$ for each read, we can write

$$p(\mathbf{Z}|\boldsymbol{\theta}) = \prod_{n=1}^{N} \prod_{m=1}^{M} \theta_m^{z_{nm}}, \tag{1}$$

where $z_{nm} \in \{0, 1\}$ is a binary variable which indicates whether the n^{th} fragment came from the m^{th} transcript ($z_{nm} = 1$) and is subject to $\sum_{m=0}^{M} z_{nm} = 1$. We use $\mathbf{Z}$ to represent the collection of all allocation vectors. We note that both $\boldsymbol{\theta}$ and $\mathbf{Z}$ are variables to be inferred, with $\boldsymbol{\theta}$ the main object of interest. $\boldsymbol{\theta}$ can be transformed later into some more convenient measure, for instance reads per kilobase of length per million sequenced reads (RPKM) (Mortazavi et al., 2008), though it is more convenient from a probabilistic point of view to work with $\boldsymbol{\theta}$ directly. The variables $\mathbf{Z}$ are sometimes known in the machine learning literature as latent variables. Although not of interest directly, inference of these variables is essential in order to infer $\boldsymbol{\theta}$.

Read model

An important part of the model is the likelihood term $p(\mathbf{r}_n|T_m)$ which is the probability of generating the n^{th} read from the m^{th} transcript. Writing the collection of all reads as $\mathbf{R} = \{\mathbf{r}_n\}_{n=1}^{N}$, the likelihood given a set of alignments $\mathbf{Z}$ is

$$p(\mathbf{R}|\mathbf{T}, \mathbf{Z}) = \prod_{n=1}^{N} \prod_{m=1}^{M} p(\mathbf{r}_n|T_m)^{z_{nm}}, \tag{2}$$

where T_m represents the m^{th} transcript and $\mathbf{T}$ represents the transcriptome. The values of $p(\mathbf{r}_n|T_m)$ for all alignments can be computed before performing inference in $\boldsymbol{\theta}$ since we are assuming a known transcriptome. For paired-end reads, the mates originate from a single fragment and their likelihood is inferred jointly. Denoting $\mathbf{r}_n = (r_n^{(1)}, r_n^{(2)})$, the likelihood of alignment is computed as

$$P(\mathbf{r}_n|T_m) = P(l|T_m)P(p|l, T_m) \prod_{i=1,2} P(r_n^{(i)}|seq_{mlp}) \ , \tag{3}$$

where l is the length of a fragment, p is its position and seq_{mlp} denotes the underlying reference sequence. The fragment length distribution can be predefined or inferred empirically. The position likelihood, $P(p|l, T_m)$, can be either uniform or account for different biases using an empirical model as in Glaus et al. (2012). The last term, $\prod_{i=1,2} P(r_n^{(i)}|seq_{mlp})$ describes the probability of observed read sequences based on quality scores and base discrepancy between read and reference. For detailed description of the alignment likelihood estimation please refer to Glaus et al. (2012).

Identifying noisy reads

Our model is similar to previous work (Glaus et al., 2012), but does not contain a variable identifying reads as belonging to a 'noise' class. To circumvent the explicit formulation of a model with this variable, we introduce a 'noise transcript' which we append to the list of known transcripts. The generative probability of any read from this transcript, $p(\mathbf{r}_n|T_0)$, is again calculated according to the model described in Glaus et al. (2012). Due to the conjugate relationships between the variables in our model and those of Glaus et al. (2012), the models are the same, subject to a slight reformulation of the prior parameters.

Prior over $\boldsymbol{\theta}$

The final part of our model is to specify some prior belief in the vector $\boldsymbol{\theta}$. To make our approximations tractable, it is necessary to use a conjugate prior, which in this case is a Dirichlet distribution

$$p(\boldsymbol{\theta}) = \frac{\Gamma(\hat{\alpha}^o)}{\prod_{m=1}^{M} \Gamma(\alpha_m^o)} \prod_{m=1}^{M} \theta_m^{\alpha_m^o - 1} \tag{4}$$

where α_m^o represents our prior belief in the values of θ_m and $\hat{\alpha}^o = \sum_{m=1}^{M} \alpha_m^o$. We use a weak but proper prior $\alpha_m^o = 1$; $m = 0 \dots M$ which corresponds to a single "pseudo-count" read (or read-pair) for each transcript.

2.2 Approximate inference

We are interested in computing the posterior distribution for the mixing proportions, $p(\boldsymbol{\theta} \,|\, \mathbf{R}, \mathbf{T}) \propto \sum_{\mathbf{Z}} p(\mathbf{R} \,|\, \mathbf{T}, \mathbf{Z}) p(\mathbf{Z} \,|\, \boldsymbol{\theta}) p(\boldsymbol{\theta})$. For very small datasets, it is possible to perform exact Bayesian inference in this model, however for any realistically sized problem, exact inference is impossible due to the combinatorial explosion of the number of possible solutions. Our proposed solution is to use a collapsed version of Variational Bayes (VB). VB involves approximating the posterior probability density of all the model parameters with another distribution q,

$$q(\boldsymbol{\theta}, \mathbf{Z}) \approx p(\boldsymbol{\theta}, \mathbf{Z}|\mathbf{R}, \mathbf{T}). \tag{5}$$

The approximation is optimised by minimising the Kullback-Leibler (KL) divergence between $q(\boldsymbol{\theta}, \mathbf{Z})$ and $p(\boldsymbol{\theta}, \mathbf{Z}|\mathbf{R}, \mathbf{T})$ (Bishop, 2006). To make the VB approach tractable, some factorisations need to be assumed in the approximate

posterior. In the case of the current model, we assume that the posterior probability of the transcript proportions factorises from the alignments:

$$q(\boldsymbol{\theta}, \mathbf{Z}) = q(\boldsymbol{\theta})q(\mathbf{Z}). \tag{6}$$

Further factorisations in $q(\mathbf{Z})$ occur due to the simplicity of the model, revealing $q(\mathbf{Z}) = \prod_{n=1}^{N} q(\mathbf{z}_n)$. We write the approximate distribution for $q(\mathbf{Z})$ using the parameters ϕ_{nm}:

$$q(\mathbf{Z}) = \prod_{n=1}^{N} \prod_{m=1}^{M} \phi_{nm}^{z_{nm}}. \tag{7}$$

We need not introduce parameters for $q(\boldsymbol{\theta})$ since it will arise implicitly in our derivation in terms of ϕ.

The objective function

Approximate inference is performed by optimisation: the parameters of the approximating distribution are changed so as to minimise the KL divergence. Whilst the KL divergence is not computable, it is possible to derive a lower bound on the marginal likelihood, maximisation of which minimises the KL divergence (see e.g. Bishop, 2006). Here we derive a lower bound which is dependent only on the parameters of $q(\mathbf{Z})$, with the optimal distribution for $q(\boldsymbol{\theta})$ arising implicitly for any given $q(\mathbf{Z})$. First we construct a lower bound on the conditional log probability of the reads $\mathbf{R}$ given the transcript proportions $\boldsymbol{\theta}$ and the known transcriptome $\mathbf{T}$:

$$\begin{aligned}
\ln p(\mathbf{R} \,|\, \mathbf{T}, \boldsymbol{\theta}) &= \ln \int p(\mathbf{R} \,|\, \mathbf{Z}, \mathbf{T}) p(\mathbf{Z} \,|\, \boldsymbol{\theta}) \,\mathrm{d}\mathbf{Z} \\
&\geq \mathbb{E}_{q(\mathbf{Z})} \Big[\ln p(\mathbf{R} \,|\, \mathbf{Z}, \mathbf{T}) + \ln p(\mathbf{Z} \,|\, \boldsymbol{\theta}) - \ln q(\mathbf{Z}) \Big] \\
&= \sum_{n=1}^{N} \sum_{m=1}^{M} \phi_{nm} \big(\ln p(\mathbf{r}_n \,|\, T_m) + \ln \theta_m - \ln \phi_{nm} \big) \\
&= \mathcal{L}_1(\boldsymbol{\theta}),
\end{aligned} \tag{8}$$

where the first line follows from Jensen's inequality in a similar fashion to standard VB methods. We have denoted this conditional bound $\mathcal{L}_1(\boldsymbol{\theta})$, which is still a function of $\boldsymbol{\theta}$. In order to generate a bound on the marginal likelihood, $p(\mathbf{R} \,|\, \mathbf{T})$, we need to remove this dependence on $\boldsymbol{\theta}$ which we do in a Bayesian fashion, by substituting $\mathcal{L}_1(\boldsymbol{\theta})$ into the following Bayesian marginalisation:

$$\begin{aligned}
p(\mathbf{R} \,|\, \mathbf{T}) &= \int p(\mathbf{R} \,|\, \mathbf{T}, \boldsymbol{\theta}) p(\boldsymbol{\theta}) \,\mathrm{d}\boldsymbol{\theta} \\
&\geq \int \exp\{\mathcal{L}_1(\boldsymbol{\theta})\} p(\boldsymbol{\theta}) \,\mathrm{d}\boldsymbol{\theta}.
\end{aligned} \tag{9}$$

Solving this integral and taking the logarithm gives us our final bound which equates to

$$\begin{aligned} \ln p(\mathbf{R} \,|\, \mathbf{T}) \geq \mathcal{L} = \sum_{n=1}^{N} \sum_{m=1}^{M} \phi_{nm} \big(\ln p(\mathbf{r}_n \,|\, T_m) - \ln \phi_{nm} \big) \\ + \ln \Gamma(\hat{\alpha}^o) - \ln \Gamma(\hat{\alpha}^o + N) - \sum_{m=1}^{M} \Big(\ln \Gamma(\alpha_m^o) - \ln \Gamma(\alpha_m^o + \hat{\phi}_m) \Big), \end{aligned} \tag{10}$$

where $\hat{\phi}_m = \sum_{n=1}^{N} \phi_{nm}$ and we also have that the approximate posterior distribution for $\boldsymbol{\theta}$ is a Dirichlet distribution with parameters $\alpha_m^o + \hat{\phi}_m$.

2.3 Optimisation

Having established the objective function as a lower bound on the marginal likelihood, all that remains is to optimise the variables of the approximating distribution $q(\mathbf{Z}, \boldsymbol{\theta})$. The dimensionality of this optimisation is rather high and potentially rather difficult. Optimisation in standard VB is usually performed by an EM like algorithm, which performs a series of convex optimisations in each of the factorised variables alternately. In our formulation of the problem, we only need to optimise the parameters of the distribution $q(\mathbf{Z})$, which we do by a gradient-based method. Taking a derivative of (10) with respect to the parameters ϕ gives

$$\frac{\partial \mathcal{L}}{\partial \phi_{nm}} = \ln p(\mathbf{r}_n \,|\, T_m) - \ln \phi_{nm} - 1 + \psi(\alpha_m^o + \hat{\phi}_m), \tag{11}$$

where ψ is the digamma function. To avoid constrained optimisation we reparameterise ϕ as γ:

$$\phi_{nm} = \frac{e^{\gamma_{nm}}}{\sum_{m'=1}^{M} e^{\gamma_{nm'}}} \tag{12}$$

and it is then possible to optimise the variables γ using a standard gradient-based optimiser.

2.4 Geometry

Information geometry concerns the interpretation of statistical objects in a geometric fashion. Specifically, a class of probability distributions behaves as a Riemannian manifold with curvature given by the Fisher information. Amari (1998) showed that the direction of the steepest descent on a such a manifold is given by the natural gradient:

$$\widetilde{\nabla} \mathcal{L} = G^{-1} \nabla \mathcal{L} \ , \tag{13}$$

where G is the Fisher information matrix. Since we are performing optimisation of the distribution $q(\mathbf{Z})$, we can make use of the natural gradient in computing a search direction (Honkela et al., 2010). For our problem, we assume that

the $N \times M$ matrix $\mathbf{Z}$ has been transformed into a NM vector, and the Fisher information corresponding to γ_{nm}, $\gamma_{n'm'}$ is given by

$$G[m,n,m',n'] = \begin{cases} \phi_{nm} - \phi_{nm}^2, & \text{if } n = n' \text{ and } m = m' \\ -\phi_{nm}\phi_{nm'}, & \text{if } n = n' \text{ but } m \neq m' \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

We note that this structure is block-diagonal, and that each block can be easily inverted using the Sherman-Morrison identity, giving an analytical expression for $G^{-1}\nabla\mathcal{L}$, and thus making the natural gradient very fast to compute (see Hensman et al. (2015) for more details). One can draw comparisons with a Newton method, where G would be replaced with a Hessian, though in the proposed case the system is much cheaper to compute.

The optimisation of the variational parameters then proceeds as follows. Following random initialisation, a unit step is taken in the natural gradient direction. Subsequent steps are subject to conjugate gradients (Honkela et al., 2010). If the conjugate gradient step should fail to improve the objective we revert to a VBEM update, which is guaranteed to improve the bound. For more details, see Hensman et al. (2012).

2.5 Truncation

The optimisation described above has $N \times M$ free parameters for optimisation, one to align each read to each transcript. However, for most read-transcript pairs, $p(\mathbf{r}_n \,|\, T_m)$ will be negligibly small. We follow Glaus et al. (2012) in truncating the values of $p(\mathbf{r}_n \,|\, T_m)$ to zero for reads which do not suitably align. Examining the objective function (10) we see that we can also set ϕ_{nm} to zero for these truncated alignments (using the convention that $0\ln(0) = 0$) and thus also $\gamma_{nm} = -\infty$ for the same. This truncation dramatically reduces the computational load of our algorithm, reducing the dimensionality of the optimisation space as well as reducing the number of operations needed to compute the objective.

2.6 The approximate posterior

Having fitted our model, we may wish to propagate the posterior distribution through a second set of processing, for example to identify differentially expressed transcripts as in BitSeq stage 2 (Glaus et al., 2012). Whilst it may be desirable to solve both stages together in a Bayesian framework, the size of the problem generally forbids this, therefore we propose the use of either a moment-matching or sampling procedure to propagate $q(\boldsymbol{\theta})$ through further analysis. The approximate posterior $q(\boldsymbol{\theta})$ is a Dirichlet distribution, whose marginals have the following useful properties:

$$\mathbb{E}\Big[\theta_m\Big] = \frac{\alpha_m^o + \hat{\phi}_m}{\hat{\alpha}^o + N}, \tag{15}$$

$$\text{var}[\theta_m] = (\alpha_m^o + \hat{\phi}_m)(\hat{\alpha}^o + N - \alpha_m^o - \hat{\phi}_m)C, \tag{16}$$

$$\text{cov}[\theta_m, \theta_m'] = -(\alpha_m^o + \hat{\phi}_m)(\alpha_{m'}^o + \hat{\phi}_{m'})C, \tag{17}$$

with $C = (\hat{\alpha}^o + N)^{-2}(\hat{\alpha}^o + N + 1)^{-1}$. This approximate posterior is somewhat inflexible, in that it cannot express arbitrary covariances between the transcripts. This arises from the factorising assumption amongst the assignment of reads to transcripts: reads are assigned independently in the variational method and their dependence cannot be modelled. This is reflected in the results section where we show empirically that the VB approximation leads to an underestimation of the variance. Nonetheless, this simplifying assumption leads to very accurate expression estimates much faster than MCMC.

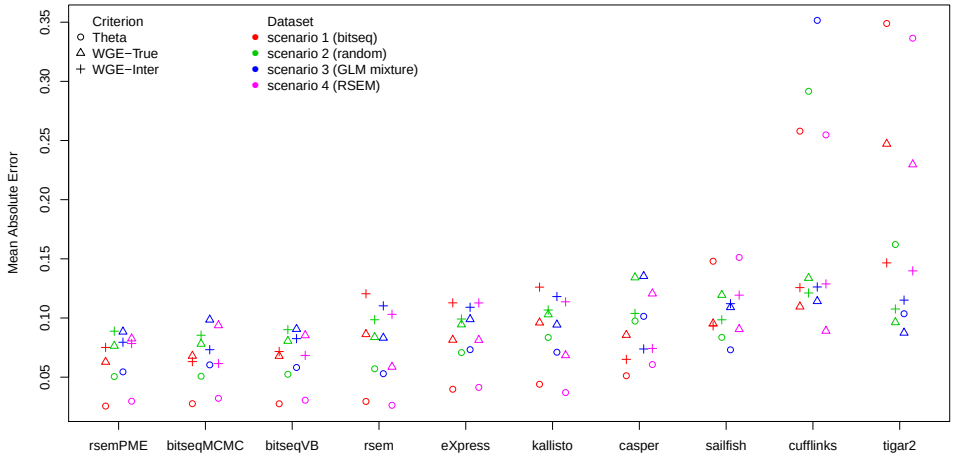

Figure 2: Ranking of methods for five replicates of simulated RNA-seq reads. WGE-Inter: inter-replicate consistency of within gene estimates, WGE-True: within gene estimates compared to the true values and Theta: estimated relative transcript expression compared to the true values. Scores have been nomrmaised to unity per dataset. Alternatives normalizations do not change the ranking (see supplementary material).

3 Results and Discussion

The proposed BitSeqVB algorithm was compared to Cufflinks (Trapnell et al., 2010), RSEM as well as the corresponding MCMC sampler RSEM-PME (Li and Dewey, 2011), BitSeqMCMC (Glaus et al., 2012), eXpress (Roberts and Pachter, 2013), Casper (Rossell et al., 2014), Sailfish (Patro et al., 2014), Tigar2 (Nariai et al., 2014) and Kallisto (Bray et al., 2015). We note that both MCMC samplers (RSEM-PME and BitSeqMCMC) use similar collapsed Gibbs algorithms but are initialised differently: RSEM-PME starts from the ML solution found by RSEM while BitSeqMCMC starts from a random initialisation and therefore requires more interations to find a good solution.

We used two main ways for benchmarking: analysis on synthetic data allowed comparison with a known ground truth under a variety of generative scenarios; analysis on high-quality replicated human data focussed on inter-replicate consistency following the evaluation of Rossell et al. (2014). We find BitSeqVB to

have excellent inter-replicate consistency and accuracy, closely approximating the original MCMC algorithm, while also being competitive with other methods in terms of run-time. We subsequently analyse in more detail the approximation to the posterior used in the BitSeqVB method. For comparison with other methods, we used default settings where appropriate: both MCMC sampling methods use 1000 posterior samples as default. However, this number refers to effective samples (Gelman et al., 2003) in BitSeqMCMC and not to single iterations as in RSEM-PME. We turned off creating of unecessary output files in RSEM. The experiments were conducted on a 4 core workstation. All the details of the experiments can be found at the aforementioned url.

3.1 Inference accuracy on synthetic data

RNA-seq reads from $M = 48009$ transcripts of the UCSC/hg19 transcriptome annotation (Kent et al., 2002) were simulated using the Spanki software (Sturgill et al., 2013). The expression is evaluated in three different measures: transcript expression accuracy (Theta), transcript within-gene relative proportion accuracy (WGE-True) and inter-replicate consistency (WGE-Inter).

A ground truth was generated using four different models of transcript expression, according to the following scenarios:

1. estimated expression levels from real data using BitSeqMCMC (≈ 56 million reads per replicate)

2. randomly selected expression levels according to a uniform distribution defined on the set $(10, 200)$ (≈ 7.8 million reads per replicate)

3. a high-dimensional mixture of Poisson Generalized Linear models, which was recently used to model the heterogeneity in RNA-seq datasets (Papastamoulis et al., 2014b) (≈ 5.5 million reads per replicate)

4. estimated expression levels from real data using RSEM (≈ 18 million reads per replicate)

For each scenario five replicates are generated according to a Negative Binomial model. Full details of the four scenarios are described in the Supplementary Material. Finally, the resulting reads-per-kilobase (RPK) values were fed into Spanki. Next, the simulated reads were aligned to the reference annotation using Bowtie2 and/or Tophat2. In particular, BitSeq, RSEM, eXpress and Tigar require transcriptomic alignments so Bowtie2 (version 2.0.6) (Langmead and Salzberg, 2012) was used, while Cufflinks and Casper work with genomic alignments using Tophat2 and Bowtie2. On the other hand, Sailfish and Kallisto produce their own alignments using k-mers mapping and pseudo-alignments, respectively. The corresponding mapping rate for genomic or transcriptomic alignments was 96%. The same amount of reads pseudo-aligned when using Kallisto, whilst Sailfish mapped a smaller portion of k-mers ($\approx 63\%$).

Figure 2 displays the mean absolute error (MAE) according to the three criteria, after performing the following normalization:

$$\sum_{m \in \text{methods}} \text{MAE}_m^{(c)} = 1,$$

$\forall c \in \{\text{Theta}, \text{WGE-Inter}, \text{WGE-True}\}$, in order to make all criteria equally weighted for each scenario. Moreover, the "Theta" and "WGE-True" metrics were averaged across the five replicates, while "WGE-Inter" was averaged across all ten combinations of pairs of replicates. The methods were ranked with respect to their average across the three criteria. RSEM-PME, BitSeqMCMC and BitSeqVB are ranked as best when considering all three criteria. RSEM has similar accuracy in terms of the ground truth expression (Theta and WGE-True) but has lower inter-replicate consistency (WGE-Inter). Conversely, Casper achieves good performance with respect to inter-replicate consistency (WGE-Inter) but is less accurate in comparison to the ground truth values (WGE-True and Theta). The ranking of methods with respect to run-time is shown in Figure 3. Note that the run-time calculation excludes the alignment procedure, but includes all other computations (including computing alignment probabilities in BitSeq's case). An exception is made for Sailfish and Kallisto, where alignment is not required, making these by far the fastest methods. Timings which include the time required for alignment are provided in Supplementary Figure 19.

The plots of inter-replicate consistency between pairs of replicates are shown in the supplementary material (Figures 2, 4, 6 and 8). As seen there, Kallisto, RSEM, Sailfish, Tigar2, Cufflinks and eXpress, produce estimates close to the boundary of the parameter space. This is also obtained for RSEM-PME except for scenario 2. This behaviour is avoided when using BitSeqMCMC, BitSeqVB and Casper.

The accuracy of BitSeqVB is very close to the two sampling methods BitSeqMCMC and RSEM-PME, but it is consistently faster that these approaches, being about 10 times faster than BitSeqMCMC and 2 times faster than RSEM-PME on average (RSEM-PME is significantly faster than BitSeqMCMC because is uses many fewer iterations of MCMC). BitSeqVB has similar speed to the eXpress method in most cases whilst exhibiting much better accuracy.

We conclude that the proposed VB algorithm is competitive in speed while exhibiting both high accuracy and good inter-replicate consistency.

3.2 Replicate consistency in human data

A recent study (Rossell et al., 2014) used the mean absolute error between pairs of replicates of the same ENCODE experiment in order to assess the accuracy of transcript expression estimation methods. For this purpose, the relative within gene expression estimates are used (WGE-Inter). Here, we provide an extended version of this analysis in order to benchmark against BitSeqMCMC and seven other methods.

In total, five ENCODE datasets (Tilgner et al., 2012) consisting of 2×76 bp reads were selected, corresponding to the following pairs of replicates:

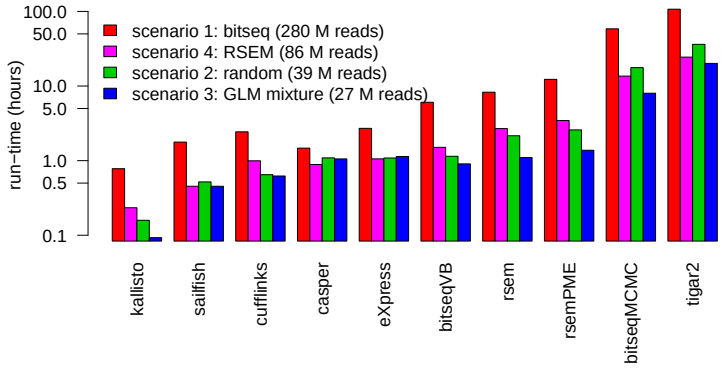

Figure 3: Run-time in hours (log-scale) for 4 synthetic data samples with 5 replicates per sample. The total number of simulated reads is shown in parenthesis.

(SRR307897, SRR307898), (SRR307901, SRR307902), (SRR307907, SRR307908), (SRR307911, SRR307912), (SRR307915, SRR307916). All methods were applied assuming the same UCSC/hg19 transcriptome annotation as in the previous section. According to the alignment rates shown in Figure 4(a), all methods work with almost the same number of mapped reads when Bowtie2 is used. This is not the case for Bowtie1 which for some reason fails on this data set.

Figure 4(b) illustrates the ranking of methods in terms of the MAE criterion, averaged across the five datasets. We conclude that BitSeqMCMC has best inter-replicate consistency, closely followed by BitSeqVB, while Casper comes next. Sailfish, RSEM, Tigar2 and Cufflinks exhibit almost two times larger MAE, while eXpress is almost 2.5 times worse according to this measure. Based on these five samples there is a partial order: BitSeqMCMC $\succ$ BitSeqVB $\succ$ {Casper, Kallisto} $\succ$ RSEM-PME $\succ$ {RSEM, Sailfish} $\succ$ {Cufflinks, Tigar2} $\succ$ eXpress, where $\succ$ denotes "is better in every experiment". Excluding BitSeq (MCMC and VB) and Casper, we see that many methods produced estimates close to the boundary of the parameter space, as seen in Figure 5. This means that many transcripts are estimated as weakly or non-expressed in one replicate while being highly expressed in the other. This problem appears to affect methods using Maximum Likelihood estimation (RSEM, Sailfish, Cufflinks, eXpress) or Bayesian methods using a very weak prior (Tigar2). Casper ensures consistency with a strong prior, but this may degrade the accuracy of absolute estimates relative to BitSeq because of stronger regularisation. We note that Casper uses MAP parameter estimation, finding the mode of the posterior distribution, while the BitSeq methods estimate the mean of the posterior distribution. Using the posterior mean may avoid spurious values where the mode is a long way from the mass of the posterior without the need for an overly strong prior. Finally, note that the coherency of inter-replicate consistency estimates in our simulation study (supplementary Figures 2,4,6 and 8) with the one reported here.

The run-time for each method is displayed in Figure 4(c). BitSeqVB is comparable to the fastest methods (except for Kallisto which is by far the fastest

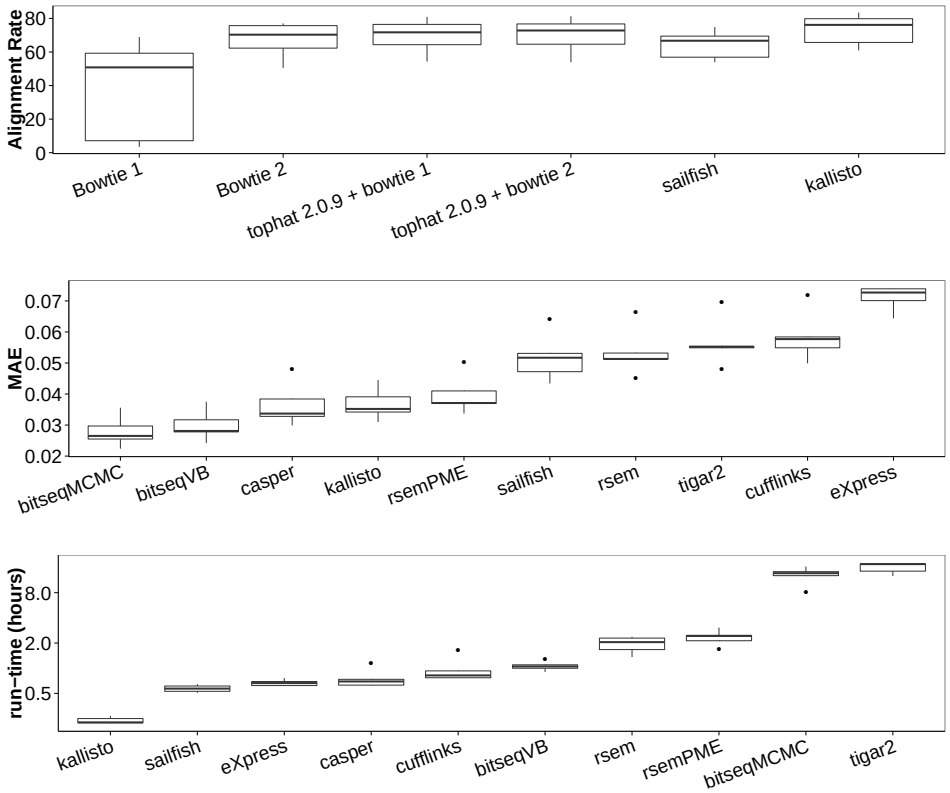

Figure 4: Five ENCODE pairs of replicates. (a): Alignment rates for transcriptome mapping (Bowtie1 and Bowtie2), genome mapping (Tophat 2.0.9 with Bowtie1 and Bowtie2), k-mers mapping (Sailfish) and pseudo-alignments (Kallisto). (b): Ranking of methods in terms of the Mean Absolute Error. (c): Run-time in hours (log-scale) with 24.6M (mapped) reads per sample.

method) while being ranked as second in terms of the MAE criterion. We conclude that BitSeqVB offers perhaps the best trade-off in accuracy and runtime on these datasets.

Finally, we mention that the BitSeqMCMC performance here is in stark contrast with the performance reported in Rossell et al. (2014). The reason for this is that in Rossell et al. (2014) reads were aligned using Bowtie1 whereas we are using Bowtie2. As seen in Figure 4(a), Bowtie1 can exhibit very low alignment rates for these samples. Interestingly, this behaviour is not present when Bowtie1 is combined with Tophat for genome mapping. The low alignment rates of Bowtie1 means that methods have available only a tiny fraction of the useful data, leading to less accurate results. This explains the weak agreement of same transcript estimates between pairs of replicates reported for BitSeq in Rossell et al. (2014) and is a reminder that it is very important to check the alignment rates.

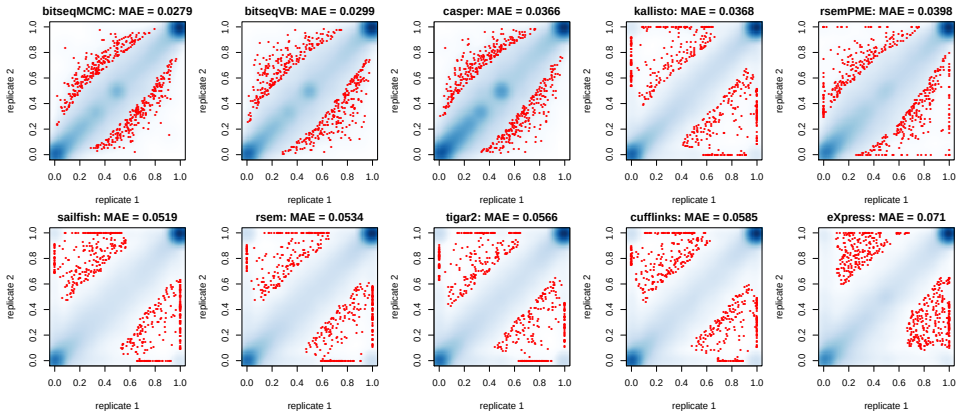

Figure 5: Scatterplots of within gene estimates for one pair of replicates (SRR307907 and SRR307908) from the ENCODE data. The blue color corresponds to a smoothed color density representation of the scatterplot.

3.3 Analysis of the Variational Bayes approximation

To examine the properties of the variational approximation, we focussed on ENCODE dataset SRR307907 (Tilgner et al., 2012). This contained 30.8 million reads, each 76bp. The reads were again mapped to the same UCSC/hg19 reference transcriptome resulting in 23.7 million mapped reads.

Our main potential concern in using the VB method is the quality of approximation to the posterior. Figure 6(a) shows a comparison of the variational posterior with a ground truth computed by MCMC with a very large sampling time. We conclude that the VB method consistently provides very accurate estimates of the posterior mean across the whole range of expression levels. The estimates of posterior variance are less consistent and for a fraction of transcripts the variances are underestimated (Figure 6(b)). It appears that VB only estimates the Poisson variance associated with random sampling of reads (Figure 6(d)), whereas the true posterior variance is larger for some transcripts due to the uncertainty in assigning multi-mapping reads (Figure 6(c)). If estimation of the expression level is all that is required, then it would seem that the VB method suffices. However, downstream methods which make use of uncertainty in the transcript quantification (such as the differential expression analysis proposed in BitSeq stage 2 (Glaus et al., 2012)) may suffer from the poor approximation in terms of posterior variance. This can potentially be addressed by augmenting the VB method with a more accurate approximation as done in a recent study that proposed a new VB algorithm with improved variance estimates and a tighter lower bound on the log-marginal likelihood (Papastamoulis et al., 2014a).

3.4 Convergence comparison

We further investigate convergence properties of MCMC and VB in terms of mean expression. RNA-seq data was obtained from ENCODE experiment SRX110318, run SRR387661, generating 124.8 million 76 bp read-pairs. We mapped the reads

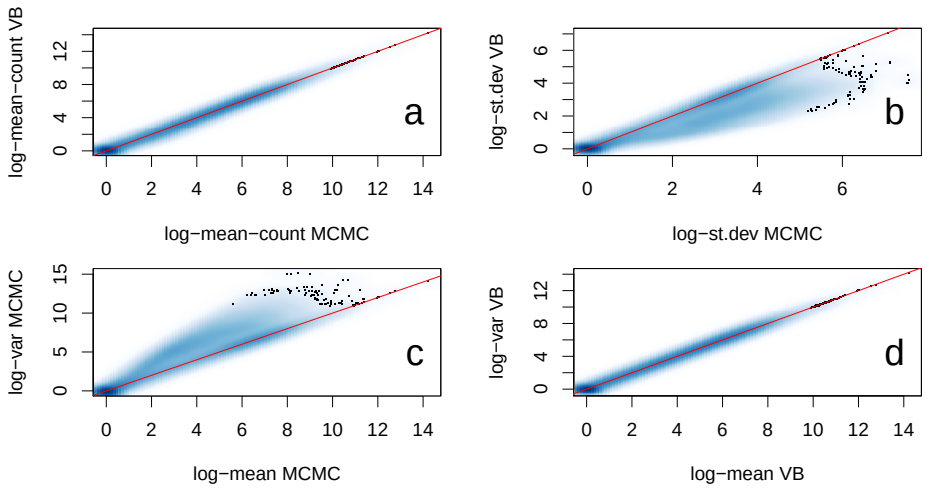

Figure 6: A comparison of the first two moments of the approximate posterior expression in counts per transcript: (a) posterior mean (R^2 correlation is 0.999) (b) posterior standard deviation: the VB method significantly under-estimates the posterior variance (σ^2). (c), (d) posterior mean-variance relationship in MCMC and VB respectively. Shading represents the number of transcripts in each region.

using Bowtie 2 to a reference transcriptome using 8713 transcripts of chromosome 19 from Ensembl human cDNA, release 70 (Flicek et al., 2013).

As the true expression levels are unknown, we used a long run of MCMC as the ground truth for mean expression estimates. Running the inference methods for a certain number of iterations, we record the run time and calculate Root Mean Square Error (RMSE) of estimated expression. The convergence of our variational method (BitSeqVB) and the original Gibbs sampling procedure (BitSeqMCMC) is shown in Figure 7. We also include a standard implementation of VB (similar to Nariai et al. (2013)) but using the BitSeq model (denoted VBEM). It is straightforward to derive this algorithm from our VB algorithm derivation since standard VBEM is obtained as a special case of steepest descent VB learning (Hensman et al., 2012). Our implementation of VB converges first in about 2 minutes. Surprisingly, some runs of collapsed MCMC converge to better estimates even faster than standard VB, which takes around 10 minutes. However, as MCMC is a stochastic method, an estimate that is consistently better than the results obtain by VB is only obtained after 900 minutes.

4 Conclusion

We have presented a new Variational Bayes method for inference of transcript expression from RNA-seq data. Building on previous work in BitSeq, we have presented a fast approximate inference method. The mean of the posterior distribution of expression levels was very well estimated in substantially less time than the original MCMC algorithm. The method is therefore suitable when

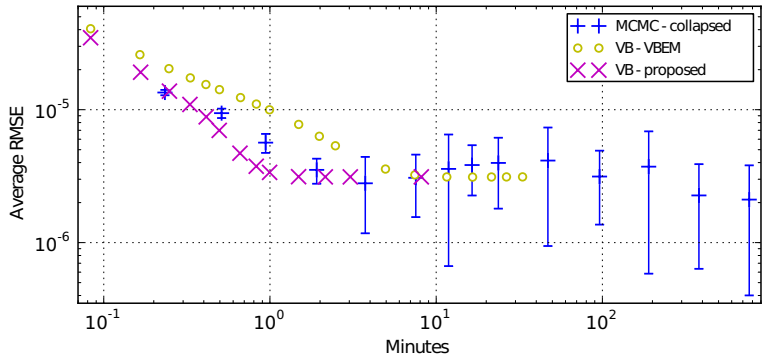

Figure 7: Convergence comparison of Collapsed MCMC with standard VB algorithm and VB with Fletcher-Reeves conjugate gradient optimisation. Expression estimates obtained by very long run of MCMC are used as a ground truth and average root mean square error over 10 runs was calculated, two standard deviations are used as error bars. The VB methods with several randomised initial conditions showed negligible differences in convergence.

point estimates of expression are sufficient, especially if time and computational resources are limited. We have compared both the original BitSeq algorithm and our new method with the majority of available methods for transcript expression estimation and conclude that BitSeqVB is highly competitive both in terms of expression estimation and run-time. We also note that an existing VBEM algorithm implementation, TIGAR, does not provide a significant improvement over Gibbs sampling in terms of computational time in our examples, as well as having a very high memory requirement.

The newest method considered here, Kallisto, is found to be extremely fast and perform with very good accuracy compared to other maximum likelihood approaches. This speed-up is achieved through avoiding full alignment and simplifying the likelihood computation through using a pseudo-alignment approach. However, the method still produces estimates at the boundary in our between-replicate comparisons similar to all maximum likelihood methods. It would therefore be very interesting to apply a Bayesian algorithm, such as the fast VB method proposed here, using the same likelihood model as Kallisto.

Finally, we suggest some areas for future development. The fast and consistent convergence of the VB method makes it useful for quick examination of the data before the Gibbs sampler is run. Further, since it provides an excellent approximation to the mean of the posterior, it could be used to e.g. reduce the burn-in time for the Gibbs sampler, or as the initial stage of a more sophisticated approximating technique, as in Papastamoulis et al. (2014a).

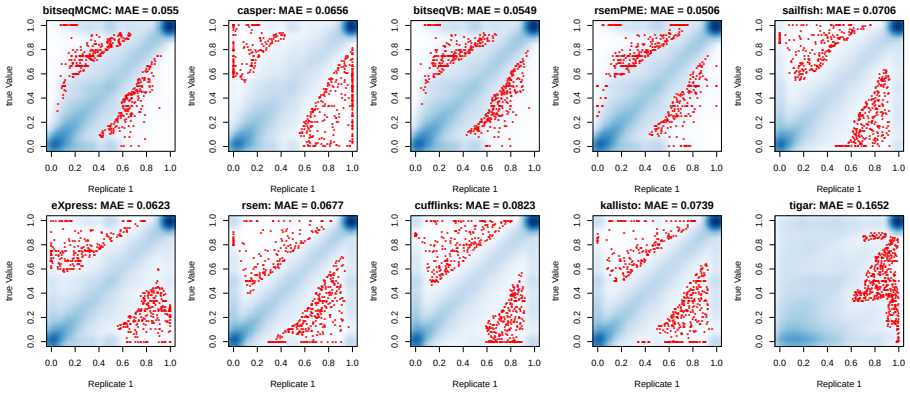

Figure 8: Scenario 1: Scatterplots of within gene estimates for one replicate of the simulated data versus the true values. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

5 Supplementary Data

We provide details of the simulation study and additional Figures as Supplementary Material which is available online.

Funding JH was supported by an MRC fellowship and BBSRC award BB/H018123/2, PP and MR by BBSRC award BB/J009415/1 and MR by EU FP7 award "RADIANT" (grant no. 305626). PG was supported by the Engineering and Physical Sciences Research Council [EP/P505208/1]. AH was supported by the Academy of Finland [259440 to AH].

Acknowledgements We thank three anonymous reviewers for useful comments which have greatly improved the paper and Lior Pachter for his blog comments on a preliminary version of this paper which led us to include more realistic data simulation scenarios.

A Simulation study details

We describe the generative process of the RPK values used in the simulation section of the manuscript. Let $\mathcal{NB}(\mu, \phi)$ denote the Negative Binomial distribution, with mean equal to μ and variance equal to $\mu + \mu^2/\phi$, $\mu \geqslant 0$, $\phi > 0$. Denote by RPK_{jm} the RPK value for transcript $m = 1, \ldots, M$ at replicate $j = 1, \ldots, J$, where J and M denote the number of replicates and the total number of transcripts, respectively.

A.1 Scenario 1: BitSeq estimates from real data

We used human data (SRR307907) from the ENCODE project in order to capture the dynamics of realistic RNA-seq datasets. BitSeq (MCMC) was used to

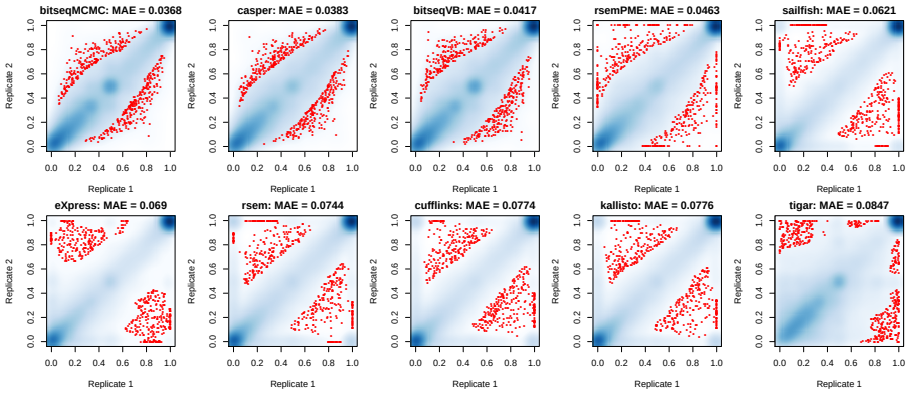

Figure 9: Scenario 1: Scatterplots of within gene estimates for one pair of replicates from the simulated RNA-seq reads. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

estimate the relative expression levels of $M = 48009$ transcripts. Next, the resulting estimates was used as input to generate the baseline mean. Reads were simulated according to the following generative process:

$$\mathrm{RPK}_{jm} \sim \mathcal{NB}(\widehat{\mu}_m, 50), \quad m = 1, \ldots, M, j = 1, \ldots, 5.$$

with $\widehat{\mu}_m$, $m = 1, \ldots, M$ denoting the corresponding estimates of RPK values according to BitSeq MCMC. This resulted in 56 million paired-end reads of 76 base-pairs per replicate, almost 280 million reads in total.

Figures 8 and 9 illustrate the scatterplots of within gene estimates for the first replicate versus the true values and the second replicate estimates, respectively. The plots are ordered according to the inter-replicate consistency Mean Absolute Error, as shown in Figure 9. Note that BitSeqMCMC, BitSeqVB and Casper are the only methods which avoid extreme outliers on the boundary of the inter-replicate consistency graphs.

A.2 Scenario 2: random selection

A ground truth was generated using randomly selected levels of transcript expression, and the resulting reads-per-kilobase (RPK) values were fed into Spanki. Let $\mathcal{U}(\alpha, \beta)$ denotes the uniform distribution defined on the set (α, β), $\alpha < \beta$. Reads were simulated according to the following generative process:

$$\begin{aligned} \mu_m &\sim \mathcal{U}(10, 200), \quad m = 1, \ldots, M \\ \mathrm{RPK}_{jm} &\sim \mathcal{NB}(\mu_m, 20), \quad m = 1, \ldots, M, j = 1, \ldots, 5. \end{aligned}$$

This resulted in 15.8 million paired-end reads of 76 base-pairs per replicate, almost 80 million reads in total.

Figures 10 and 11 illustrate the scatterplots of within gene estimates for the first replicate versus the true values and the second replicate estimates, respectively. The plots are ordered according to the inter-replicate consistency Mean

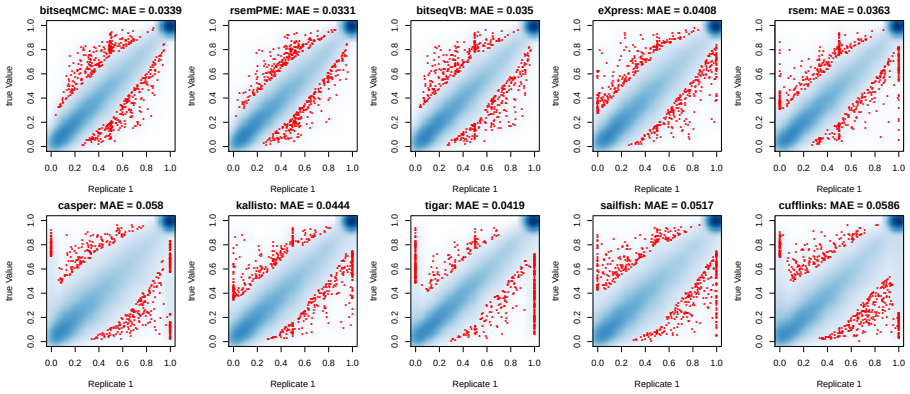

Figure 10: Scenario 2: Scatterplots of within gene estimates for one replicate of the simulated data versus the true values. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

Absolute Error, as shown in Figure 11. Note that RSEM-PME, BitSeqMCMC and BitSeqVB are the only methods which avoid extreme outliers on the boundary of these graphs.

A.3 Scenario 3: mixture of Poisson GLMs

A ground truth was generated using a mixture of Poisson Generalized Linear models. It has been recently demostrated that this approach can be used in order to model the underlying heterogeneity in RNA-seq datasets (Papastamoulis et al., 2014b). Let $\mathcal{P}(\lambda)$ denotes the Poisson distribution with mean $\lambda > 0$ and let also $\mathcal{D}(\alpha_1, \ldots, \alpha_K)$ denotes the Dirichlet distribution with $\alpha_k > 0$, $k = 1, \ldots, K$ for a given positive integer K. Reads were simulated according to the following generative process:

$$\begin{aligned} K &= 20 \\ (\pi_1, \ldots, \pi_K) &\sim \mathcal{D}(1, \ldots, 1) \\ x_m &\sim \mathcal{U}(0, 30) \\ \log \lambda_{jkm} &= \alpha_k + \beta_{jk} x_m \\ (\mu_{1m}, \ldots, \mu_{5m}) &\sim \sum_{k=1}^{K} \pi_k \prod_{j=1}^{5} \mathcal{P}(\lambda_{jkm}), \\ \mathrm{RPK}_{jm} &\sim \mathcal{NB}(\mu_{jm}, 50), \end{aligned}$$

where $m = 1, \ldots, M$, $j = 1, \ldots, 5$ and $k = 1, \ldots, K$. The regression coefficients (α_k, β_{jk}) are generated as follows.

$$\begin{aligned} \alpha_k &= \frac{K^* - k}{2}, \quad k = 1, \ldots, K, \\ \beta_{jk} &= \begin{cases} \frac{10 - \alpha_k}{10}, & k = 1, \ldots, K,\ j = 1 \\ \beta_{1k} + \varepsilon_{jk}, & j \geqslant 2,\ \varepsilon_{jk} \sim \mathcal{N}(0, 0.001^2) \end{cases} \end{aligned}$$

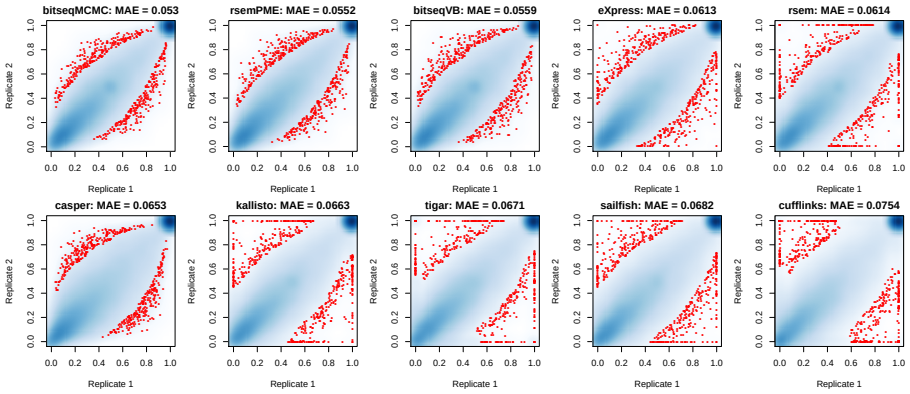

Figure 11: Scenario 2: Scatterplots of within gene estimates for one pair of replicates from the simulated RNA-seq reads. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

where $K^* := \lfloor 0.5 + (K+1)/2 \rfloor$ ($\lfloor \alpha \rfloor$ stands for the integer part of α).

This resulted in 5.5 million paired-end reads of 76 base-pairs per replicate, almost 27.5 million reads in total. Figures 12 and 13 illustrate the scatterplots of within gene estimates for the first replicate versus the true values and the second replicate estimates, respectively. The plots are ordered according to the inter-replicate consistency Mean Absolute Error, as shown in Figure 13. Note that BitSeqMCMC, BitSeqVB and Casper are the only methods which avoid extreme outliers on the boundary of the inter-replicate consistency graphs.

A.4 Scenario 4: RSEM estimates from real data

We used human data (replicates SRR307907 and SRR307908) from the ENCODE project in order to capture the dynamics of realistic RNA-seq datasets. RSEM was used to estimate the relative expression levels of $M = 48009$ transcripts for each dataset. Next, the resulting estimates were averaged and used as input to generate the baseline mean. Reads were simulated according to the following generative process:

$$\mathrm{RPK}_{jm} \quad \sim \quad \mathcal{NB}(\widehat{\mu}_m, 50), \quad m = 1, \ldots, M, j = 1, \ldots, 5.$$

with $\widehat{\mu}_m$, $m = 1, \ldots, M$ denoting the corresponding estimates of RPK values according to RSEM. This resulted in 18 million paired-end reads of 76 base-pairs per replicate, almost 86 million reads in total.

Figures 14 and 15 illustrate the scatterplots of within gene estimates for the first replicate versus the true values and the second replicate estimates, respectively. The plots are ordered according to the inter-replicate consistency Mean Absolute Error, as shown in Figure 15. Note that BitSeqMCMC, BitSeqVB and Casper are the only methods which avoid extreme outliers on the boundary of the inter-replicate consistency graphs.

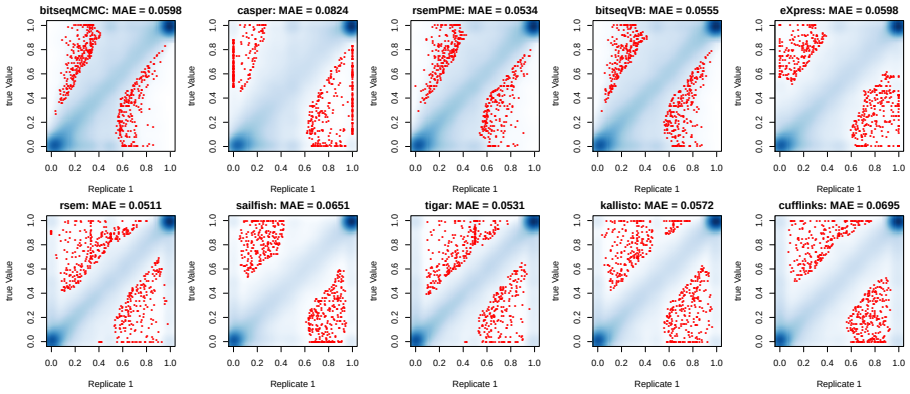

Figure 12: Scenario 3: Scatterplots of within gene estimates for one replicate of the simulated data versus the true values. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

The overall ranking of methods is shown in Figure 16. Now we used an alternative normalisation by dividing each score by the minimum per criterion, so that the normalised score of the best method is equal to 1.

B Evaluation measures

The following measures are used.

Theta: $\frac{1}{J}\sum_{j=1}^{J}\frac{1}{M}\sum_{m=1}^{M}|\log\widehat{\theta}_m^{(j)} - \log\theta_m^{(j)}|$, where $\theta_m^{(j)}$ and $\widehat{\theta}_m^{(j)}$ denote true and estimated relative transcript expression for transcript m of replicate j, $m = 1,\ldots,M$, $j = 1,\ldots,J$.

WGE-True: $\frac{1}{J}\sum_{j=1}^{J}\frac{1}{M}\sum_{m=1}^{M}|\widehat{q}_m^{(j)} - q_m^{(j)}|$, where $q_m^{(j)} := \frac{\theta_m^{(j)}}{\sum_{k\in T_m}\theta_m^{(j)}}$ and $\widehat{q}_m^{(j)} := \frac{\widehat{\theta}_m^{(j)}}{\sum_{k\in T_m}\widehat{\theta}_m^{(j)}}$ denote the true and estimated relative within gene transcript expression levels and also $T_m := \{k :$ transcript k in same gene as transcript m, $k = 1,\ldots,M\}$, $m = 1,\ldots,M$ denotes the set of transcripts belonging to the parent gene of each transcript.

WGE-Inter: $\sum_{i<j}\frac{2}{J(J-1)}\sum_{m=1}^{M}\frac{1}{M}\sum_{m=1}^{M}|\widehat{q}_m^{(i)} - \widehat{q}_m^{(j)}|$.

In addition we provide in Figure 17 the ranking of methods when the within gene estimates are computed according to the TPM measure, instead of θ. Figure 18 displays a comparison of ranking of methods when using only a subset of highly expressed transcripts.

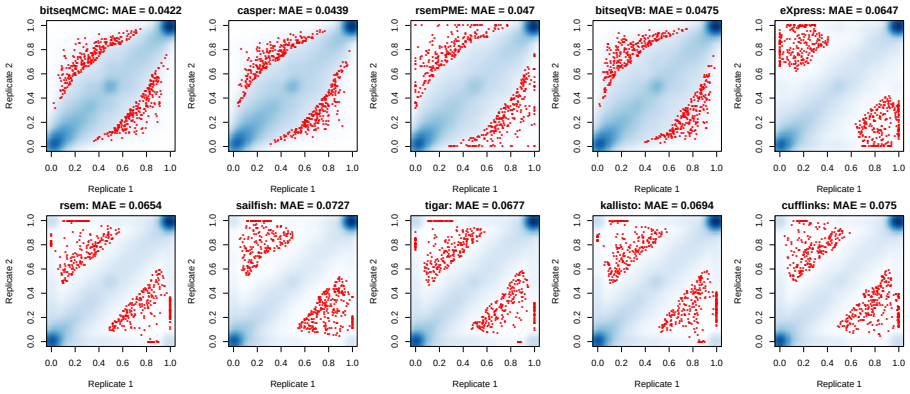

Figure 13: Scenario 3: Scatterplots of within gene estimates for one pair of replicates from the simulated RNA-seq reads. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

References

Amari, S. (1998). Natural gradient works efficiently in learning. Neural computation, 10(2):251–276.

Bishop, C. (2006). Pattern recognition and machine learning. Springer New York, New York.

Bray, N., Pimentel, H., Melsted, P., and Pachter, L. (2015). Near-optimal RNA-Seq quantification. arXiv (q-bio.QM), arXiv:1505.02710v2.

Flicek, P. et al. (2013). Ensembl 2013. Nucleic acids research, 41(Database issue):D48–55.

Gelman, A., Carlin, J., Stern, H., and Rubin, D. (2003). Bayesian Data Analysis, Second Edition. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.

Glaus, P., Honkela, A., and Rattray, M. (2012). Identifying differentially expressed transcripts from RNA-seq data with biological variation. Bioinformatics, 28(13):1721–1728.

Hensman, J., Rattray, M., and Lawrence, N. D. (2012). Fast variational inference in the conjugate exponential family. In Advances in Neural Information Processing Systems (NIPS).

Hensman, J., Rattray, M., and Lawrence, N. D. (2015). Fast nonparametric clustering of structured time-series. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 37(2):383–393.

Honkela, A., Raiko, T., Kuusela, M., Tornio, M., and Karhunen, J. (2010). Approximate Riemannian conjugate gradient learning for fixed-form variational Bayes. Journal of Machine Learning Research, 11:3235–3268.

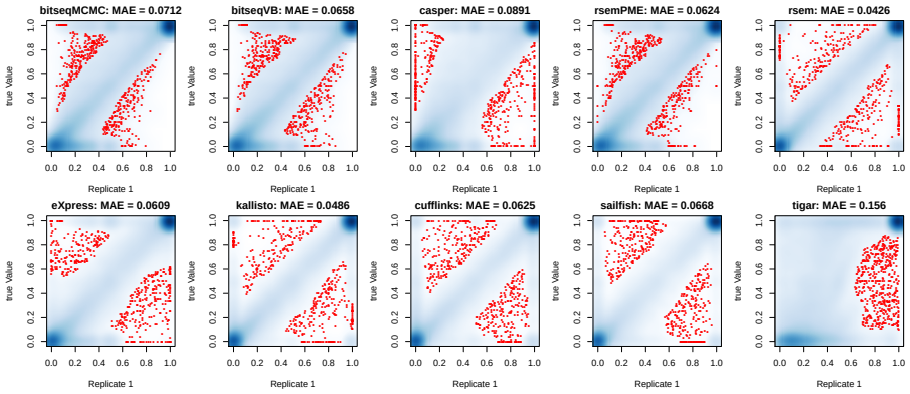

Figure 14: Scenario 4: Scatterplots of within gene estimates for one replicate of the simulated data versus the true values. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

Jiang, H. and Wong, W. H. (2009). Statistical inferences for isoform expression in rna-seq. Bioinformatics, 25(8):1026–1032.

Katz, Y., Wang, E. T., Airoldi, E. M., and Burge, C. B. (2010). Analysis and design of RNA sequencing experiments for identifying isoform regulation. Nat Methods, 7(12):1009–1015.

Kent, W., Sugnet, C., Furey, T., Roskin, K., Pringle, T., Zahler, A., and Haussler, D. (2002). The human genome browser at UCSC. Genome Research, 6(12):996–1006.

Langmead, B. and Salzberg, S. L. (2012). Fast gapped-read alignment with Bowtie 2. Nature methods, 9(4):357–9.

Li, B. and Dewey, C. N. (2011). RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. BMC Bioinformatics, 12:323.

Li, B., Ruotti, V., Stewart, R. M., Thomson, J. A., and Dewey, C. N. (2010). RNA-Seq gene expression estimation with read mapping uncertainty. Bioinformatics, 26(4):493–500.

Mortazavi, A., Williams, B. A., McCue, K., Schaeffer, L., and Wold, B. (2008). Mapping and quantifying mammalian transcriptomes by RNA-Seq. Nature Methods, 5(7):621–8.

Nariai, N., Hirose, O., Kojima, K., and Nagasaki, M. (2013). TIGAR: transcript isoform abundance estimation method with gapped alignment of RNA-Seq data by variational Bayesian inference. Bioinformatics, 18(29):2292–2299.

Nariai, N., Kojima, K., Mimori, T., Sato, Y., Kawai, Yosuke Yamaguchi-Kabata, Y., and Nagasaki, M. (2014). TIGAR2: sensitive and accurate estimation of

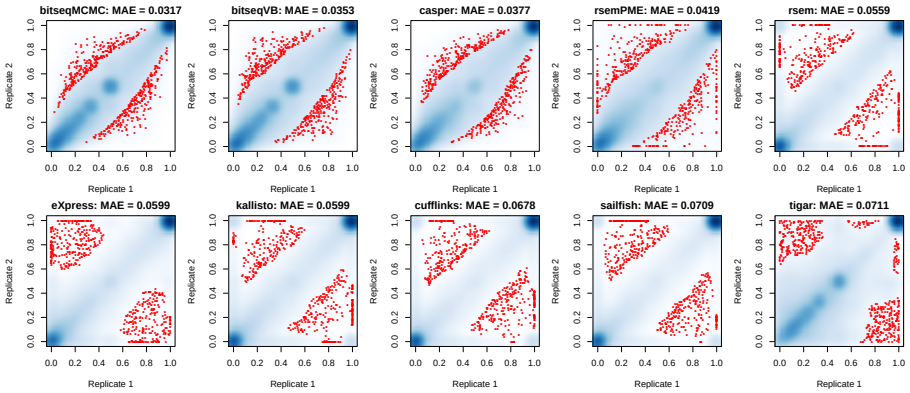

Figure 15: Scenario 4: Scatterplots of within gene estimates for one pair of replicates from the simulated RNA-seq reads. The blue color corresponds to a smoothed color density representation of the scatterplot and the red color emphasizes points from the lowest regional densities.

transcript isoform expression with longer RNA-Seq reads. BMC Genomics, (accepted).

Papastamoulis, P., Hensman, J., Glaus, P., and Rattray, M. (2014a). Improved variational Bayes inference for transcript expression estimation. Statistical Applications in Genetics and Molecular Biology, 13(2):203–216.

Papastamoulis, P., Martin-Magniette, M.-L., and Maugis-Rabusseau, C. (2014b). On the estimation of mixtures of Poisson regression models with large number of components. Computational Statistics and Data Analysis, (doi: http://dx.doi.org/10.1016/j.csda.2014.07.005).

Patro, R., Mount, S. M., and Kingsford, C. (2014). Sailfish enables alignment-free isoform quantification from RNA-seq reads using lightweight algorithms. Nature biotechnology, 32(5):462–464.

Roberts, A. and Pachter, L. (2013). Streaming fragment assignment for real-time analysis of sequencing experiments. Nature methods, 10(1):71–73.

Rossell, D., Stephan-Otto Attolini, C., Kroiss, M., and Stöcker, A. (2014). Quantifying alternative splicing from paired-end RNA-sequencing data. The Annals of Applied Statistics, 8(1):309–330.

SEQC/MAQC-III Consortium (2014). A comprehensive assessment of rna-seq accuracy, reproducibility and information content by the sequencing quality control consortium. Nat Biotechnol, 32(9):903–914.

Sturgill, J., Malone, J. H., Sun, X., Smith, H. E., Rabinow, L., Samson, M. L., and Oliver, B. (2013). Design of RNA splicing analysis null models for post hoc filtering of Drosophila head RNA-Seq data with the splicing analysis kit (Spanki). BMC Bioinformatics, 14:320.

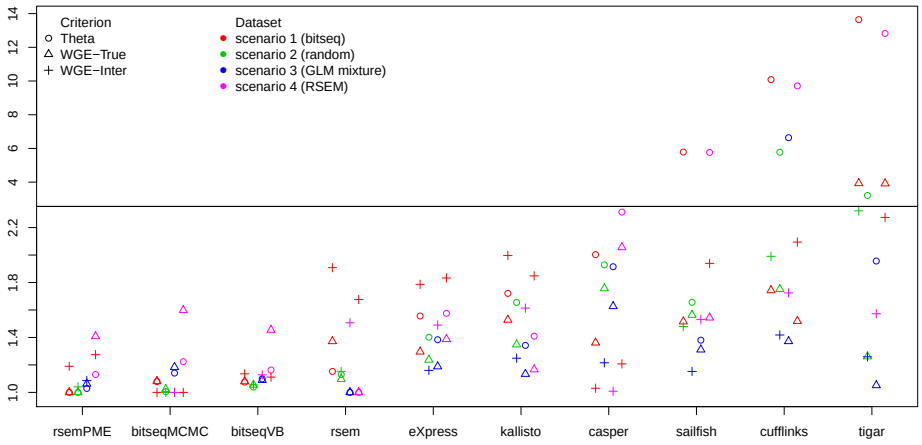

Figure 16: Ranking of methods for five replicates of simulated RNA-seq reads. WGE-Inter: inter-replicate consistency of within gene estimates, WGE-True: within gene estimates compared to the true values and Theta: estimated relative transcript expression compared to the true values. Scores have been normalised by dividing by the min score per criterion.

Tilgner, H., Knowles, D. G., Johnson, R., Davis, C. A., Chakrabortty, S., Djebali, S., Curado, J., Snyder, M., Gingeras, T. R., and Guigó, R. (2012). Deep sequencing of subcellular RNA fractions shows splicing to be predominantly co-transcriptional in the human genome but inefficient for lncRNAs. Genome Res, 22(9):1616–1625.

Trapnell, C., Hendrickson, D. G., Sauvageau, M., Goff, L., Rinn, J. L., and Pachter, L. (2013). Differential analysis of gene regulation at transcript resolution with RNA-seq. Nat Biotechnol, 31(1):46–53.

Trapnell, C., Williams, B. A., Pertea, G., Mortazavi, A., Kwan, G., van Baren, M. J., Salzberg, S. L., Wold, B. J., and Pachter, L. (2010). Transcript assembly and quantification by RNA-Seq reveals unannotated transcripts and isoform switching during cell differentiation. Nature Biotechnology, 28(5):516–520.

Turro, E., Su, S.-Y., Gonçalves, Â., Coin, L. J. M., Richardson, S., and Lewin, A. (2011). Haplotype and isoform specific expression estimation using multi-mapping RNA-seq reads. Genome Biol, 12(2):R13.

Xing, Y., Yu, T., Wu, Y. N., Roy, M., Kim, J., and Lee, C. (2006). An expectation-maximization algorithm for probabilistic reconstructions of full-length isoforms from splice graphs. Nucleic Acids Res, 34(10):3150–3160.

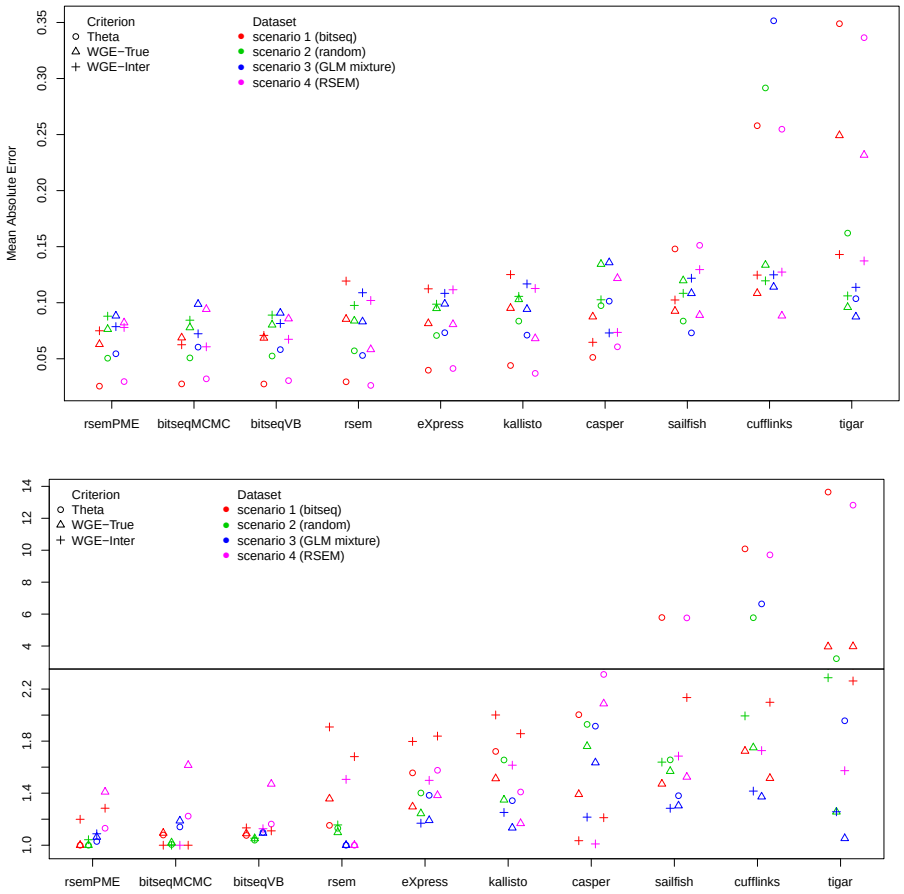

Figure 17: Ranking of methods when using the TPM instead of theta for the within gene estimates (up: normalization according to unity sum, down: normalization by dividing by the min).

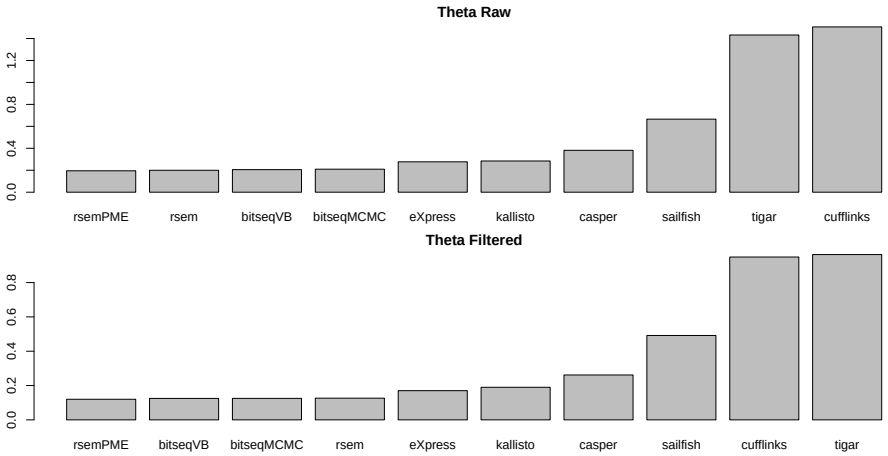

Figure 18: Ranking of methods according to Theta when using all expressed transcripts (up) and only the ones with at least 50 reads per replicate (down).

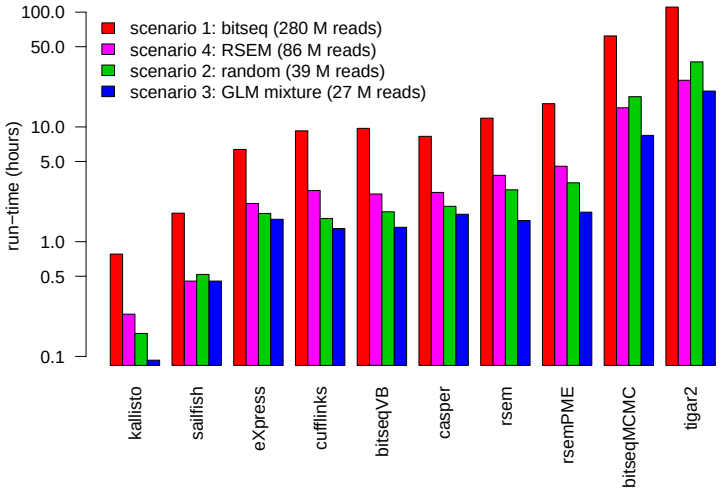

Figure 19: Runtime per method including the time needed for the alignment procedure.